Line-distortion, Bandwidth and Path-length of a graph*

Feodor F. Dragan[1], Ekkehard Köhler[2], and Arne Leitert[1]

[1] Algorithmic Research Laboratory, Department of Computer Science,
Kent State University, Kent, OH 44242, USA
dragan@cs.kent.edu, aleitert@cs.kent.edu
[2] Mathematisches Institut, Brandenburgische Technische Universität Cottbus,
D-03013 Cottbus, Germany
ekoehler@math.tu-cottbus.de

arXiv:1409.8389v1 [cs.DS] 30 Sep 2014

Abstract. For a graph $G = (V, E)$ the *minimum line-distortion problem* asks for the minimum k such that there is a mapping f of the vertices into points of the line such that for each pair of vertices x, y the distance on the line $|f(x) - f(y)|$ can be bounded by the term $d_G(x, y) \leq |f(x) - f(y)| \leq k\, d_G(x, y)$, where $d_G(x, y)$ is the distance in the graph. The *minimum bandwidth problem* minimizes the term $\max_{uv \in E} |f(u) - f(v)|$, where f is a mapping of the vertices of G into the integers $\{1, \ldots, n\}$.
We investigate the minimum line-distortion and the minimum bandwidth problems on unweighted graphs and their relations with the *minimum length* of a Robertson-Seymour's path-decomposition. The *length* of a path-decomposition of a graph is the largest diameter of a bag in the decomposition. The *path-length* of a graph is the minimum length over all its path-decompositions. In particular, we show:

- if a graph G can be embedded into the line with distortion k, then G admits a Robertson-Seymour's path-decomposition with bags of diameter at most k in G;
- for every class of graphs with path-length bounded by a constant, there exist an efficient constant-factor approximation algorithm for the minimum line-distortion problem and an efficient constant-factor approximation algorithm for the minimum bandwidth problem;
- there is an efficient 2-approximation algorithm for computing the path-length of an arbitrary graph;
- AT-free graphs and some intersection families of graphs have path-length at most 2;
- for AT-free graphs, there exist a linear time 8-approximation algorithm for the minimum line-distortion problem and a linear time 4-approximation algorithm for the minimum bandwidth problem.

Keywords: *graph algorithms; approximation algorithms; minimum line-distortion; minimum bandwidth; Robertson-Seymour's path-decomposition; path-length; AT-free graphs.*

1 Introduction and previous work

Computing a minimum distortion embedding of a given n-vertex graph G into the line ℓ was recently identified as a fundamental algorithmic problem with important applications in various areas of computer science, like computer vision [43], as well as in computational chemistry and biology (see [28,29]). It asks, for a given graph $G = (V, E)$, to find a mapping f of vertices V of G into points of ℓ with minimum number k such that $d_G(x, y) \leq |f(x) - f(y)| \leq k\, d_G(x, y)$ for every $x, y \in V$. The parameter k is called the *minimum line-distortion* of G and denoted by $\mathsf{ld}(G)$. The embedding f is called *non-contractive* since $d_G(x, y) \leq |f(x) - f(y)|$ for every $x, y \in V$.

In [4], Bădoiu et al. showed that this problem is hard to approximate within a constant factor. They gave an exponential-time exact algorithm and a polynomial-time $\mathcal{O}(n^{1/2})$-approximation algorithm for arbitrary unweighted input graphs, along with a polynomial-time $\mathcal{O}(n^{1/3})$-approximation algorithm for unweighted trees. In another paper [3] Bădoiu et al. showed that the problem is hard to approximate by a factor $\mathcal{O}(n^{1/12})$, even for weighted trees. They also gave a better polynomial-time approximation algorithm for general weighted graphs, along with a polynomial-time algorithm

* Results of this paper were partially presented at the SWAT 2014 conference [12].

that approximates the minimum line-distortion k embedding of a weighted tree by a factor that is polynomial in k.

Fast exponential-time exact algorithms for computing the line-distortion of a graph were proposed in [17,18]. Fomin et al. in [18] showed that a minimum distortion embedding of an unweighted graph into the line can be found in time $5^{n+o(n)}$. Fellows et al. in [17] gave an $\mathcal{O}(nk^4(2k+1)^{2k})$ time algorithm that for an unweighted graph G and integer k either constructs an embedding of G into the line with distortion at most k, or concludes that no such embedding exists. They extended their approach also to weighted graphs obtaining an $\mathcal{O}(nk^{4W}(2k+1)^{2kW})$ time algorithm, where W is the largest edge weight. Thus, the problem of minimum distortion embedding of a given n-vertex graph G into the line ℓ is Fixed Parameter Tractable.

Recently, Heggernes et al. in [26,27] initiated the study of minimum distortion embeddings into the line of specific graph classes. In particular, they gave polynomial-time algorithms for the problem on bipartite permutation graphs and on threshold graphs [27]. Furthermore, in [26], Heggernes et al. showed that the problem of computing a minimum distortion embedding of a given graph into the line remains NP-hard even when the input graph is restricted to a bipartite, cobipartite, or split graph, implying that it is NP-hard also on chordal, cocomparability, and AT-free graphs. They also gave polynomial-time constant-factor approximation algorithms for split and cocomparability graphs.

Table 1 and Table 2 summarise the results mentioned above.

Table 1. Existing solutions for calculating the minimum line-distortion λ.

Graph Class	Solution Quality	Run Time	Source
trees (unweighted)	$\mathcal{O}(n^{1/3})$-approx.	polynomial	[4]
trees (weighted)	$\lambda^{\mathcal{O}(1)}$-approx.	polynomial	[3]
general (unweighted)	$\mathcal{O}(n^{1/2})$-approx.	polynomial	[4]
	optimal	$5^{n+o(n)}$	[18]
	optimal	$\mathcal{O}(n\lambda^4(2\lambda+1)^{2\lambda})$	[17]
bipartite permutation	optimal	$\mathcal{O}(n^2)$	[27]
threshold	optimal	linear	[27]
split	6-approx.	linear	[26]
cocomparability	6-approx.	$\mathcal{O}(n\log^2 n+m)$	[26]

Table 2. Existing hardness results for calculating the minimum line-distortion.

Graph Class	Result	Source
general	$\mathcal{O}(1)$-approximation is NP-hard	[4]
trees (weighted)	Hard to $\mathcal{O}(n^{1/12})$-approximate	[3]
bipartite	NP-hard	[26]
cobipartite	NP-hard	[26]
split	NP-hard	[26]
AT-free	NP-hard	[26]
cocomparability	NP-hard	[26]
chordal	NP-hard	[26]

The minimum distortion embedding into the line may appear to be closely related to the widely known and extensively studied graph parameter *bandwidth*, denoted by $\mathsf{bw}(G)$. The only difference between the two parameters is that a minimum distortion embedding has to be *non-contractive*,

meaning that the distance in the embedding between two vertices of the input graph has to be at least their original distance, whereas there is no such restriction for bandwidth.

Formally, given an unweighted graph $G = (V, E)$ on n vertices, consider a 1-1 map f of the vertices V into integers in $\{1, \ldots, n\}$; f is called a *layout* of G. The *bandwidth of layout* f is defined as the maximum stretch of any edge, i.e., $\mathsf{bw}(f) = \max_{uv \in E} |f(u) - f(v)|$. The *bandwidth* of a graph is defined as the minimum possible bandwidth achievable by any 1-1 map (layout) $V \to \{1, \ldots, n\}$. That is, $\mathsf{bw}(G) = \min_{f:\, V \to \{1,\ldots,n\}} \mathsf{bw}(f)$.

It is known that $\mathsf{bw}(G) \leq \mathsf{ld}(G)$ for every connected graph G (see, e.g., [27]). However, the bandwidth and the minimum line-distortion of a graph can be very different. For example, it is common knowledge that a cycle of length n has bandwidth 2, whereas its minimum line-distortion is exactly $n-1$ [27]. Bandwidth is known to be one of the hardest graph problems; it is NP-hard even for very simple graphs like *caterpillars of hair-length at most* 3 (i.e., trees in which all the vertices are within distance 3 of a central path and all vertices of degree at least 3 are on the path) [36], and it is hard to approximate by a constant factor even for trees [2] and caterpillars with arbitrary hair-lengths [13]. Polynomial-time algorithms for the exact computation of bandwidth are known for very few graph classes, including bipartite permutation graphs [25] and interval graphs [30,33,42]. Constant-factor approximation algorithms are known for AT-free graphs [31] and convex bipartite graphs [41]. Recently, in [21] Golovach et al. showed also that the bandwidth minimization problem is Fixed Parameter Tractable on AT-free graphs by presenting an $n2^{\mathcal{O}(k \log k)}$ time algorithm. For general (unweighted) n-vertex graphs, the minimum bandwidth can be approximated within a factor of $\mathcal{O}(\log^{3.5} n)$ [15]. For n-vertex trees and chordal graphs, the minimum bandwidth can be approximated within a factor of $\mathcal{O}(\log^{2.5} n)$ [24]. For n-vertex caterpillars with arbitrary hair-lengths, the minimum bandwidth can be approximated to within a factor of $\mathcal{O}(\log n / \log \log n)$ [16].

Table 3 and Table 4 summarise the results mentioned above.

Table 3. Existing solutions for calculating the minimum bandwidth k.

Graph Class	Solution Quality	Run Time	Source
caterpillars with hair-length 1 or 2	optimal	$\mathcal{O}(n \log n)$	[1]
caterpillars with arbitrary hair-lengths	$\mathcal{O}(\log n / \log \log n)$-approx.	polynomial	[16]
general	$\mathcal{O}(\log^{3.5} n)$-approx.	polynomial	[15]
chordal	$\mathcal{O}(\log^{2.5} n)$-approx.	polynomial	[24]
AT-free	2-approx.	$\mathcal{O}(nm)$	[31]
	4-approx.	$\mathcal{O}(m + n \log n)$	[31]
	optimal	$n2^{\mathcal{O}(k \log k)}$	[21]
convex bipartite	2-approx.	$\mathcal{O}(n \log^2 n)$	[41]
	4-approx.	$\mathcal{O}(n)$	[41]
bipartite permutation	optimal	$\mathcal{O}(n^4 \log n)$	[25]
interval	optimal	$\mathcal{O}(n \log^2 n)$	[42]

Table 4. Existing hardness results for calculating the minimum bandwidth.

Graph Class	Result	Source
trees	hard to approximate by a constant factor	[2]
caterpillars with arbitrary hair-lengths	hard to approximate by a constant factor	[13]
caterpillars with hair-length at most 3	NP-hard	[36]
convex bipartite	NP-hard	[41]

Our main tool in this paper is Robertson-Seymour's path-decomposition and its length. A *path-decomposition* [40] of a graph $G = (V, E)$ is a sequence of subsets $\{X_i : i \in I\}$ ($I := \{1, 2, \ldots, q\}$) of vertices of G, called *bags*, with three properties:

1. $\bigcup_{i \in I} X_i = V$;
2. For each edge $uv \in E$, there is a bag X_i such that $u, v \in X_i$;
3. For every three indices $i \leq j \leq k$, $X_i \cap X_k \subseteq X_j$. Equivalently, the subsets containing any particular vertex form a contiguous subsequence of the whole sequence.

We denote a path-decomposition $\{X_i : i \in I\}$ of a graph G by $\mathcal{P}(G)$.

The *width* of a path-decomposition $\mathcal{P}(G) = \{X_i : i \in I\}$ is $\max_{i \in I} |X_i| - 1$. The *path-width* of a graph G, denoted by $\mathsf{pw}(G)$, is the minimum width over all path-decompositions $\mathcal{P}(G)$ of G [40]. The caterpillars with hair-length at most 1 are exactly the graphs with path-width 1 [38].

Following [10] (where the notion of tree-length of a graph was introduced), we define the *length* of a path-decomposition $\mathcal{P}(G)$ of a graph G to be $\lambda := \max_{i \in I} \max_{u,v \in X_i} d_G(u, v)$ (i.e., each bag X_i has diameter at most λ in G). The *path-length* of G, denoted by $\mathsf{pl}(G)$, is the minimum length over all path-decompositions of G. Interval graphs (i.e., the intersection graphs of intervals on a line) are exactly the graphs with path-length 1; it is known (see, e.g., [9,19,20,22]) that G is an interval graph if and only if G has a path-decomposition with each bag being a maximal clique of G.

Note that these two graph parameters (path-width and path-length) are not related to each other. For instance, a clique on n vertices has path-length 1 and path-width $n - 1$, whereas a cycle on $2n$ vertices has path-width 2 and path-length n.

Following [11], where the notion of tree-breadth of a graph was introduced, we define the breadth of a path-decomposition as follows. Let $D_G(v_i, r)$ be the disk of radius r around vertex v_i, more precisely, $D_G(v_i, r) = \{w \in V : d_G(v_i, w) \leq r\}$. Then the *breadth* of a path-decomposition $\mathcal{P}(G)$ of a graph G is the minimum integer r such that for every $i \in I$ there is a vertex $v_i \in V$ with $X_i \subseteq D_G(v_i, r)$ (i.e., each bag X_i can be covered by a disk $D_G(v_i, r)$ of radius at most r in G). Note that vertex v_i does not need to belong to X_i. The *path-breadth* of G, denoted by $\mathsf{pb}(G)$, is the minimum breadth over all path-decompositions of G. Evidently, for any graph G with at least one edge, $1 \leq \mathsf{pb}(G) \leq \mathsf{pl}(G) \leq 2\,\mathsf{pb}(G)$ holds. Hence, if one parameter is bounded by a constant for a graph G then the other parameter is bounded for G as well.

Recently, Robertson-Seymour's tree-decompositions with bags of bounded radius proved to be very useful in designing an efficient approximation algorithm for the problem of minimum stretch embedding of an unweighted graph in to its spanning tree [11]. The decision version of the problem is the *tree t-spanner problem* which asks, for a given graph $G = (V, E)$ and an integer t, whether a spanning tree T of G exists such that $d_T(x, y) \leq t\, d_G(x, y)$ for every $x, y \in V$. It was shown in [11] that:

(1) if a graph G can be embedded into a spanning tree with stretch t, then G admits a Robertson-Seymour tree-decomposition with bags of radius at most $\lceil t/2 \rceil$ and diameter at most t in G (i.e., the tree-breadth $\mathsf{tb}(G)$ of G is at most $\lceil t/2 \rceil$ and the tree-length $\mathsf{tl}(G)$ of G is at most t);
(2) there is an efficient algorithm which constructs for an n-vertex unweighted graph G with $\mathsf{tb}(G) \leq \rho$ a spanning tree with stretch at most $2\rho \log_2 n$.

As a consequence, an efficient $(\log_2 n)$-approximation algorithm was obtained for embedding an unweighted graph with minimum stretch into its spanning tree [11].

1.1 Contribution of this paper

Motivated by [11], in this paper, we investigate possible connections between the line-distortion and the path-length (path-breadth) of a graph. We show that for every graph G, $\mathsf{pl}(G) \leq \mathsf{ld}(G)$

and $\mathsf{pb}(G) \leq \lceil \mathsf{ld}(G)/2 \rceil$ hold. Furthermore, we demonstrate that for every class of graphs with path-length bounded by a constant, there is an efficient constant-factor approximation algorithm for the minimum line-distortion problem. As a consequence, every graph G with $\mathsf{ld}(G) = c$ can be embedded in polynomial time into the line with distortion at most $\mathcal{O}(c^2)$ (reproducing a result from [4]). Additionally, using the same technique, we show that, for every class of graphs with path-length bounded by a constant, there is an efficient constant-factor approximation algorithm for the minimum bandwidth problem.

We also investigate (i) which particular graph classes have constant bounds on path-length and (ii) how fast the path-length of an arbitrary graph can be computed or sharply estimated. We present an efficient 2-approximation (3-approximation) algorithm for computing the path-length (resp., the path-breadth) of a graph. We show that AT-free graphs and some well-known intersection families of graphs have small path-length and path-breadth. In particular, the path-breadth of every permutation graph and every trapezoid graph is 1 and the path-length (and therefore, the path-breadth) of every cocomparability graph and every AT-free graph is at most 2. Using this and some additional structural properties, we give a linear time 8-approximation algorithm for the minimum line-distortion problem and a linear time 4-approximation algorithm for the minimum bandwidth problem for AT-free graphs.

As a consequence of our results we obtain also that convex bipartite graphs and caterpillars with hairs of bounded length admit constant factor approximations of the minimum bandwidth and the minimum line-distortion. Furthermore, the minimum line-distortion problem and the minimum bandwidth problem, are both NP-hard on bounded path-length graphs.

2 Preliminaries and metric properties of graphs with bounded path-length

All graphs occurring in this paper are connected, finite, unweighted, undirected, loopless and without multiple edges. We call $G = (V, E)$ an *n-vertex m-edge graph* if $|V| = n$ and $|E| = m$. In this paper we consider only graphs with $n > 1$. A *clique* is a set of pairwise adjacent vertices of G. By $G[S]$ we denote the subgraph of G induced by the vertices of $S \subseteq V$. By $G \setminus S$ we denote the subgraph of G induced by the vertices $V \setminus S$, i.e., the graph $G[V \setminus S]$. For a vertex v of G, the sets $N_G(v) = \{w \in V : vw \in E\}$ and $N_G[v] = N_G(v) \cup \{v\}$ are called the *open neighborhood* and the *closed neighborhood* of v, respectively.

In a graph G the *length* of a path from a vertex v to a vertex u is the number of edges in the path. The *distance* $d_G(u, v)$ between vertices u and v is the length of a shortest path connecting u and v in G. For a set $S \subseteq V$ the *diameter* of S in G is $\max_{x,y \in S} d_G(x, y)$ and its *radius* in G is $\min_{x \in V} \max_{y \in S} d_G(x, y)$ (in some papers these terms are called the *weak diameter* and the *weak radius* to indicate that the distances are measured in G not in $G[S]$). The distance between a vertex v and a set S of G is given by $d_G(v, S) = \min_{u \in S} d_G(v, u)$. The *disk* of radius k centered at vertex v in G is the set of all vertices at distance at most k from v, i.e., $D_G(v, k) = \{w \in V : d_G(v, w) \leq k\}$.

The following result generalizes a characteristic property of the famous class of AT-free graphs (see [7]). An independent set of three vertices such that each pair is joined by a path that avoids the neighborhood of the third is called an *asteroidal triple*. A graph G is an *AT-free graph* if it does not contain any asteroidal triples [7].

Proposition 1. *Let G be a graph with $\mathsf{pl}(G) \leq \lambda$. Then, for every three vertices u, v, w of G there is one vertex, say v, such that the disk of radius λ centered at v intercepts every path connecting u and w, i.e., after the removal of the disk $D_G(v, \lambda)$ from G, u and w are not in the same connected component of $G \setminus D_G(v, \lambda)$.*

Proof. Consider a path-decomposition $\mathcal{P}(G) = \{X_i : i \in I\}$ of G with length $\mathsf{pl}(G) \leq \lambda$. Consider any three vertices u, v, w of G. If any two of them, say u and v, belong to same bag X_i of $\mathcal{P}(G)$ then the disk $D_G(v, \lambda)$ contains vertex u and hence intercepts every path of G connecting vertices u and w. Assume now, without loss of generality, that all bags containing vertex u have smaller indexes in I than all bags containing vertex v, and, in turn, all bags containing vertex v have smaller indexes in I than all bags containing vertex w. Then, by properties of path-decompositions (see [9,40]), every u, w-path of G contains at least one vertex in each of the bags between the bags containing u and the bags containing w in the sequence $\{X_i : i \in I\}$. Hence, every bag X_i containing v intercepts every path connecting u and w in G. Since X_i is a subset of $D_G(v, \lambda)$, the proof is complete. □

Since for every graph G, $\mathsf{pl}(G) \leq 2\,\mathsf{pb}(G)$, the following statement is also true.

Corollary 1. *Let G be a graph with $\mathsf{pb}(G) \leq \rho$. Then, for every three vertices of G, the disk of radius 2ρ centered at one of them intercepts every path connecting the other two vertices.*

We will also need the following property of graphs with path-length λ. A path P of a graph G is called *k-dominating path* of G if every vertex v of G is at distance at most k from a vertex of P, i.e., $d_G(v, P) \leq k$. A pair of vertices x, y of G is called a *k-dominating pair* if every path between x and y is a k-dominating path of G. It is known that every AT-free graph has a 1-dominating pair [7].

Corollary 2. *Every graph G with $\mathsf{pl}(G) \leq \lambda$ has a λ-dominating pair.*

Proof. Consider a path-decomposition $\mathcal{P}(G) = \{X_1, X_2, \ldots, X_q\}$ of length $\mathsf{pl}(G) \leq \lambda$ of G. Consider any two vertices $x \in X_1$ and $y \in X_q$ and a path P between them in G. Necessarily, by properties of path-decompositions (see [9,40]), every path of G connecting vertices x and y has a vertex in every bag of $\mathcal{P}(G)$. Hence, as each vertex v of G belongs to some bag X_i of $\mathcal{P}(G)$, there is a vertex $u \in P$ with $u \in X_i$ and thus $d_G(v, u) \leq \lambda$. □

It is easy to see that a pair of vertices x and y is a k-dominating pair if and only if, for every vertex $w \in V \setminus (D_G(x, k) \cup D_G(y, k))$, the disk $D_G(w, k)$ separates x and y. Hence, a k-dominating pair, with minimum k, of an arbitrary graph $G = (V, E)$ with n vertices and m edges can be found in $\mathcal{O}(n^3 \log n)$ time as follows. As an outer loop use a binary search to find the minimum k. Inside this loop determine for the corresponding k whether there is a k-dominating pair by the following method. For each vertex v of G determine the connected components of $G \setminus D_G(v, k)$, label each vertex x in $G \setminus D_G(v, k)$ with its connected component, and put all these labels in a $n \times n$ matrix M, such that in $M(v, x)$ is the label of the connected component of vertex x in $G \setminus D_G(v, k)$. This matrix M can easily be determined in $\mathcal{O}(n(n + m))$ (for each vertex v remove $D_G(v, k)$ and determine the connected components of the remaining graph by a Breadth-First-Search). Now a pair of vertices x, y is a k-dominating pair if and only if the columns of x and y in M have different labels in every row corresponding to a vertex w with $w \in V \setminus (D_G(x, k) \cup D_G(y, k))$. Thus, an easy $\mathcal{O}(n^3)$ algorithm for checking whether there is a k-dominating pair in G is simply comparing for each pair of vertices the corresponding columns.

It is not very likely that there is a linear time algorithm to find a dominating pair, if it exists, since it is shown in [32] that finding a dominating pair is essentially as hard as finding a triangle in a graph. Yet, since path-decompositions with small length are closely related to k-domination one can search for k-dominating pairs in dependence of the path-length of a graph. We do not know how to find *in linear time* for an arbitrary graph G a k-dominating pair with $k \leq \mathsf{pl}(G)$. However, we can prove the following weaker result which will be useful in later sections.

Proposition 2. *Let G be an arbitrary graph for which the path-length is not necessarily known. There is a linear time algorithm that determines a k-dominating pair of G such that $k \leq 2\,\mathsf{pl}(G)$.*

Proof. Let $\mathsf{pl}(G) = \lambda$. Consider a path-decomposition $\mathcal{P}(G) = \{X_1, X_2, \ldots, X_q\}$ of G of length λ. Consider an arbitrary vertex s of G and, using a Breadth-First-Search $BFS(s, G)$ of G started at s, find a vertex x of maximum distance from s. Use a second Breadth-First-Search $BFS(x, G)$ of G that is started at x to find a vertex y of maximum distance from x. We claim that x, y is a $2\,\lambda$-dominating pair of G.

If there is a bag in $\mathcal{P}(G)$ containing both s and x, then $d_G(s, x) \leq \lambda$ and, by the choice of x, each vertex of G is within distance at most λ from s and, hence, within distance at most $2\,\lambda$ from x. Evidently, in this case, x, y is a $2\,\lambda$-dominating pair of G.

Assume now, without loss of generality, that $x \in X_i$ and $s \in X_l$ with $i < l$. Consider an arbitrary vertex v of G that belongs to only bags with indexes smaller than i. We show that $d_G(x, v) \leq 2\,\lambda$. As X_i separates v from s, a shortest path $P(s, v)$ of G between s and v must have a vertex u in X_i. We have $d_G(s, x) \geq d_G(s, v) = d_G(s, u) + d_G(u, v)$ and, by the triangle inequality, $d_G(s, x) \leq d_G(s, u) + d_G(u, x)$. Hence, $d_G(u, v) \leq d_G(u, x)$ and, since both u and x belong to same bag X_i, $d_G(u, x) \leq \lambda$. That is, $d_G(x, v) \leq d_G(x, u) + d_G(u, v) \leq 2d_G(u, x) \leq 2\,\lambda$.

If $d_G(x, y) \leq 2\,\lambda$ then, by the choice of y, each vertex of G is within distance at most $2\,\lambda$ from x and, hence, x, y is a $2\,\lambda$-dominating pair of G. So, assume that $d_G(x, y) > 2\,\lambda$, i.e., every bag of $\mathcal{P}(G)$ that contains y has index greater than i. Consider a bag X_j containing y. We have $i < j$. Repeating the arguments of the previous paragraph, we can show that $d_G(y, v) \leq 2\,\lambda$ for every vertex v that belongs to bags with indexes greater than j.

Consider now an arbitrary path P of G connecting vertices x and y. By properties of path-decompositions (see [9,40]), P has a vertex in every bag X_h of $\mathcal{P}(G)$ with $i \leq h \leq j$. Hence, for each vertex v of G that belongs to a bag X_h $(i \leq h \leq j)$, there is a vertex $u \in P$ (in that bag X_h) such that $d_G(v, u) \leq \lambda$. As $d_G(v, x) \leq 2\,\lambda$ for each vertex v from $X_{i'}$ with $i' < i$ and $d_G(v, y) \leq 2\,\lambda$ for each vertex v from $X_{j'}$ with $j' > j$, we conclude that P is a $2\,\lambda$-dominating path of G. □

In Algorithm 1, we formalize the method in the proof above to calculate a k-dominating shortest path with $k \leq 2\mathsf{pl}(G)$ in linear time.

Algorithm 1: Finding a k-dominating shortest path of G with $k \leq 2\mathsf{pl}(G)$.

Input: A graph G.
Output: A k-dominating shortest path with $k \leq 2\mathsf{pl}(G)$.

1 Select an arbitrary vertex s.
2 Find a vertex x for which the distance to s is maximal.
3 Find a vertex y for which the distance to x is maximal.
4 Output a shortest path from x to y.

The following proposition further strengthens the connections between small path-length graphs and AT-free graphs. Recall that the k-power of a graph $G = (V, E)$ is a graph $G^k = (V, E')$ such that for every $x, y \in V$ $(x \neq y)$, $xy \in E'$ if and only if $d_G(x, y) \leq k$.

Proposition 3. *For a graph G with $\mathsf{pl}(G) \leq \lambda$, $G^{2\,\lambda-1}$ is an AT-free graph.*

Proof. Let $\mathcal{P}(G)$ be a path-decomposition of length $\mathsf{pl}(G) \leq \lambda$ of G and let $G^{2\,\lambda-1} = (V, E')$ be the $(2\,\lambda - 1)$-power of G. Consider three arbitrary distinct vertices a, b and c of G. If for two of those vertices there is a bag $B \in \mathcal{P}(G)$ containing both, then a, b and c cannot be an asteroidal triple in $G^{2\,\lambda-1}$, since they do not form an independent set in $G^{2\,\lambda-1}$. Now assume that no bag contains more than one of a, b, and c. Without loss of generality, we can assume that b is in a

bag B_b between the bags containing a and the bags containing c. Assume that there is a path from a to c in $G^{2\lambda-1}$ avoiding B_b. Then, there is an edge $uv \in E' \setminus E$ such that the bags containing u and the bags containing v are separated by B_b in G. Since $uv \in E'$, there is a path of length at most $2\lambda - 1$ from u to v in G, and again, by properties of path-decompositions (see [9,40]), this path must contain a vertex $w \in B_b$. Without loss of generality, let $d_G(u,w) \leq d_G(v,w)$. Since $d_G(u,w) \leq d_G(u,v)/2 \leq \lambda - 1$, $d_G(b,u) \leq d_G(b,w) + d_G(u,w) \leq 2\lambda - 1$, i.e., $u \in D_G(b, 2\lambda - 1)$. Thus, each path from a to c of $G^{2\lambda-1}$ intersects $D_G(b, 2\lambda - 1)$, implying that a, b, c cannot form an asteroidal triple in $G^{2\lambda-1}$. □

Corollary 3. *If $\mathsf{pb}(G) \leq \rho$, then $G^{4\rho-1}$ is AT-free.*

A subset of vertices of a graph is called *connected* if the subgraph induced by those vertices is connected. We say that two connected sets S_1, S_2 of a graph G *see each other* if they have a common vertex or there is an edge in G with one end in S_1 and the other end in S_2. A family of connected subsets of G is called a *bramble* if every two sets of the family see each other. We say that a bramble $\mathcal{F} = \{S_1, \ldots, S_h\}$ of G is k-dominated by a vertex v of G if in every set S_i of $\mathcal{F}$ there is a vertex $u_i \in S_i$ with $d_G(v, u_i) \leq k$.

Proposition 4. *For a graph G with $\mathsf{pb}(G) \leq \rho$, every bramble of G is ρ-dominated by a vertex.*

Proof. Let $\mathcal{P}(G) = \{X_1, X_2, \ldots, X_q\}$ be a path-decomposition of breadth $\mathsf{pb}(G) \leq \rho$ of G. Consider an arbitrary connected set S of G. We claim that the bags of $\mathcal{P}(G)$ containing vertices of S form a continuous subsequence $\mathcal{I}(S)$ in $\{X_1, X_2, \ldots, X_q\}$. Assume, by induction on the cardinality of the set S, that the statement is true for the connected set $S' := S \setminus \{v\}$, where v is some vertex of S (obviously, there is always such a vertex in S). Let $\mathcal{I}(S')$ be the corresponding subsequence. Since S is connected, there must exist a vertex u in S' such that $uv \in E(G)$. By properties 2 and 3 of the path-decomposition (see definition on page 8), all bags containing vertex v form a continuous subsequence $\mathcal{I}(v)$ in $\{X_1, X_2, \ldots, X_q\}$ and there is a bag in $\mathcal{P}(G)$ which contains both vertices u and v. Then, necessarily, all bags containing vertices of S form a continuous subsequence in $\{X_1, X_2, \ldots, X_q\}$; it is the union of the two continuous subsequences $\mathcal{I}(S')$ and $\mathcal{I}(v)$ sharing a common bag.

Now let $\mathcal{F} = \{S_1, \ldots, S_h\}$ be an arbitrary bramble of G. For every set S_i, the bags of $\mathcal{P}(G)$ containing vertices of S_i form a continuous subsequence $\mathcal{I}(S_i)$ in $\{X_1, X_2, \ldots, X_q\}$. Since each two sets of $\mathcal{F}$ see each other, there must exist a bag in $\mathcal{P}(G)$ that contains a vertex from S_i and a vertex from S_j. So, for every $i, j \in \{1, \ldots, h\}$, the subsequences $\mathcal{I}(S_i)$ and $\mathcal{I}(S_j)$ overlap at least on one bag. By the Helly property for intervals of a line (i.e., every family of pairwise intersecting intervals has a common intersection), there must exist a bag B in $\mathcal{P}(G)$ which has a vertex from each set S_i ($i \in \{1, \ldots, h\}$). Let v be a vertex of G such that $B \subseteq D_G(v, \rho)$. Then, v necessarily ρ-dominates the bramble $\mathcal{F}$. □

The following result can be viewed as an analog of the classical Helly property for disks.

Corollary 4. *Let G be a graph with $\mathsf{pb}(G) \leq \rho$, let S be a subset of vertices of G, and let $r : S \to \mathbf{N}$ be a radius function defined on S such that the disks of the family $\mathcal{F} = \{D_G(x, r(x)) : x \in S\}$ pairwise intersect. Then the disks $\{D_G(x, r(x) + \rho) : x \in S\}$ have a nonempty common intersection.*

Proof. Since the family $\mathcal{F} = \{D_G(x, r(x)) : x \in S\}$ is a bramble of G, a vertex v of G ρ-dominating the bramble $\mathcal{F}$ belongs to all disks $\{D_G(x, r(x) + \rho) : x \in S\}$. □

3 Bandwidth of graphs with bounded path-length

In this section we show that there is an efficient algorithm that for any graph G with $\mathsf{pl}(G) = \lambda$ produces a layout f with bandwidth at most $\mathcal{O}(\lambda)\mathsf{bw}(G)$. Moreover, this statement is true even for all graphs with λ-dominating shortest paths. Recall that a shortest path P of a graph G is a *k-dominating shortest path* of G if every vertex v of G is at distance at most k from a vertex of P, i.e., $d_G(v, P) \leq k$.

We will need the following standard "local density" lemma.

Lemma 1 ([39]). *For each vertex $v \in V$ of an arbitrary graph G and each positive integer r,*

$$\frac{|D_G(v, r)| - 1}{2r} \leq \mathsf{bw}(G).$$

The main result of this section is the following.

Proposition 5. *Every graph G with a k-dominating shortest path has a layout f with bandwidth at most $(4k + 2)\,\mathsf{bw}(G)$. If a k-dominating shortest path of G is given in advance, then such a layout f can be found in linear time.*

Proof. Let $P = (x_0, x_1, \ldots, x_i, \ldots, x_j, \ldots, x_q)$ be a k-dominating shortest path of G. Consider a Breadth-First-Search-tree T_P of G started from path P, i.e., T_P is the $BFS(P, G)$-tree of G. For each vertex x_i of P, let X_i be the set of vertices of G that are located in the branch of T_P that is rooted at x_i (see Fig. 1(a) for an illustration). We have $x_i \in X_i$. Since P k-dominates G, we have $d_G(v, x_i) \leq k$ for every $i \in \{1, \ldots, q\}$ and every $v \in X_i$. Now create a layout f of G by placing all the vertices of X_i before all vertices of X_j, if $i < j$, and by placing the vertices within each X_i in an arbitrary order (see Fig. 1(b) for an illustration).

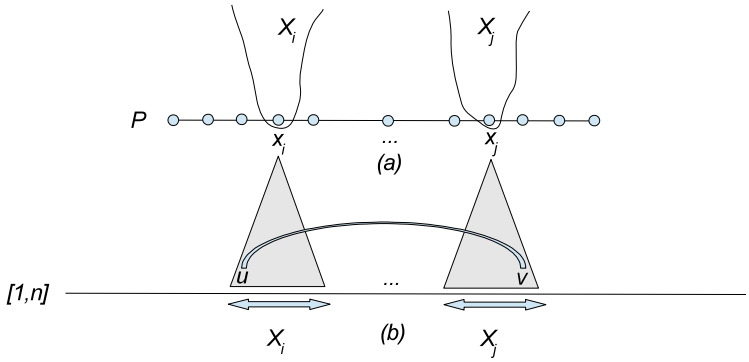

Fig. 1. Illustration to the proof of Proposition 5.

We claim that this layout f has bandwidth at most $(4k + 2)\,\mathsf{bw}(G)$. Consider any edge uv of G and assume $u \in X_i$ and $v \in X_j$ $(i \leq j)$. For this edge uv we have $f(v) - f(u) \leq |\bigcup_{\ell=i}^{j} X_\ell| - 1$. We also know that $d_P(x_i, x_j) = j - i \leq 2k + 1$, since P is a shortest path of G and $d_P(x_i, x_j) = d_G(x_i, x_j) \leq d_G(x_i, u) + 1 + d_G(x_j, v) \leq 2k + 1$. Consider vertex x_c of P with $c = i + \lfloor (j - i)/2 \rfloor$, i.e., a middle vertex of the subpath of P between x_i and x_j. Consider an arbitrary vertex w in X_ℓ, $i \leq \ell \leq j$. Since $d_G(x_c, w) \leq d_G(x_c, x_\ell) + d_G(x_\ell, w)$, $d_G(x_c, x_\ell) \leq \lceil 2k + 1 \rceil / 2$ and $d_G(x_\ell, w) \leq k$, we get $d_G(x_c, w) \leq 2k + 1$. In other words, disk $D_G(x_c, 2k + 1)$ contains all vertices of $\bigcup_{\ell=i}^{j} X_\ell$. Applying Lemma 1 to $|D_G(x_c, 2k + 1)| \geq |\bigcup_{\ell=i}^{j} X_\ell|$, we conclude $f(v) - f(u) \leq |\bigcup_{\ell=i}^{j} X_\ell| - 1 \leq |D_G(x_c, 2k + 1)| - 1 \leq 2(2k + 1)\,\mathsf{bw}(G) = (4k + 2)\,\mathsf{bw}(G)$. □

Proposition 5, Corollary 2, and Proposition 2 imply.

Corollary 5. *For every n-vertex m-edge graph G, a layout with bandwidth at most $(4\,\mathsf{pl}(G) + 2)\,\mathsf{bw}(G)$ can be found in $\mathcal{O}(n^2m)$ time and a layout with bandwidth at most $(8\,\mathsf{pl}(G)+2)\,\mathsf{bw}(G)$ can be found in $\mathcal{O}(n+m)$ time.*

Proof. For an n-vertex m-edge graph G, a k-dominating shortest path with $k \leq \mathsf{pl}(G)$ can be found in $\mathcal{O}(n^2m)$ time in the following way (see Algorithm 2). Iterate over all vertex pairs of G. For each vertex pair x, y pick a shortest x,y-path P and run $BFS(P, G)$ to find a most distant vertex v_P from P. Finally, report that path P for which $d_G(v_P, P)$ is minimum. By Corollary 2, this minimum is at most $\mathsf{pl}(G)$.

Algorithm 2: Finding a k-dominating shortest path of G with $k \leq \mathsf{pl}(G)$.

Input: A graph G.
Output: A k-dominating shortest path with $k \leq \mathsf{pl}(G)$.

1 **foreach** *vertex pair* x, y **do**
2 Find a shortest path P_{xy} from x to y.
3 Determine $k(x, y) := \max_{v\in V} d_G(v, P_{xy})$.
4 Output a path P_{xy} for which $k(x, y)$ is minimal.

Alternatively, one can use the proof of Proposition 2 to find in linear time a $2\mathsf{pl}(G)$-dominating pair x, y of G. Then, any shortest path of G between x and y is a $2\,\mathsf{pl}(G)$-dominating path of G (see Algorithm 1).

The entire method for computing a required layout is given in Algorithm 3. Its runtime and approximation ratio depend on the algorithm to calculate a k-dominating shortest path. □

Algorithm 3: An $\mathcal{O}(k)$-approximation algorithm for computing the minimum bandwidth of a graph using a k-dominating shortest path.

Input: A graph $G = (V, E)$.
Output: A layout f.

1 Find a k-dominating shortest path $P = (x_0, x_1, \ldots, x_q)$ using Algorithm 1 or Algorithm 2.
2 Partition V into sets $X_0, X_1, \ldots, X_q$ using a $BFS(P, G)$-tree of G (see the proof of Proposition 5).
3 Create a layout f of G by placing all the vertices of X_i before all vertices of X_j, if $i < j$, and by placing vertices within each X_i in an arbitrary order.
4 Output f.

Thus, we have the following interesting conclusion.

Theorem 1. *For every class of graphs with path-length bounded by a constant, there is an efficient constant-factor approximation algorithm for the minimum bandwidth problem.*

The above results did not require a path-decomposition of length $\mathsf{pl}(G)$ of a graph G as input; we also avoided the construction of such a path-decomposition of G and just relied on the existence

of a k-dominating shortest path in G. If, however, a path-decomposition with length λ of a graph G is given in advance together with G, then a better approximation ratio for the minimum bandwidth problem on G can be achieved.

Proposition 6. *If a graph G is given together with a path-decomposition of G of length λ, then a layout f with bandwidth at most $\lambda\,\mathsf{bw}(G)$ can be found in $\mathcal{O}(n^2 + n\log^2 n)$ time.*

Proof. Let $\mathcal{P}(G) = \{X_i : i \in I\}$ be a path-decomposition of length λ of $G = (V, E)$. We form a new graph $G^+ = (V, E^+)$ from G by adding an edge between a pair of vertices $u, v \in V$ if and only if u and v belong to a common bag in $\mathcal{P}(G)$. From this construction, we conclude that G is a subgraph of G^+ and G^+ is a subgraph of G^λ. It is a well-known fact (see, e.g., [5,9,19,22]) that G^+ is an interval graph and $\mathcal{P}(G) = \{X_i : i \in I\}$ gives a path-decomposition of G^+ (with $\{X_i : i \in I\}$ being cliques of G^+). In [42], an $\mathcal{O}(n\log^2 n)$ time algorithm to compute a minimum bandwidth layout of an n-vertex interval graph is given. Let f be an optimal layout produced by that algorithm for our interval graph G^+. We claim that this layout f, when considered for G, has bandwidth at most $\lambda\,\mathsf{bw}(G)$. Indeed, following [31], we have $\max_{uv\in E}|f(u) - f(v)| \le \max_{uv\in E^+}|f(u) - f(v)| = \mathsf{bw}(G^+) \le \mathsf{bw}(G^\lambda) \le \lambda\,\mathsf{bw}(G)$. Clearly, raising a graph to the λth power can only increase its bandwidth by a factor of λ. □

We formalize the method described above in Algorithm 4.

Algorithm 4: A λ-approximation algorithm for computing the minimum bandwidth for a graph with path-length λ.

Input: A graph G with a path-decomposition $\mathcal{P}(G) = \{X_1, \ldots, X_q\}$.
Output: A layout f.

1. Create a new graph $G^+ = (V, E^+)$ by adding an edge between each pair of vertices $u, v \in V$ if and only if u and v belong to a common bag in $\mathcal{P}(G)$.
2. Compute the minimum bandwidth layout f of the interval graph G^+ by using an optimal $\mathcal{O}(n\log^2 n)$ time algorithm from [42].
3. Output f.

We do not know how hard it is for an arbitrary graph to construct its path-decomposition with minimum length. We suspect that this is an NP-hard problem as the problem to check whether a graph has tree-length at most λ is NP-complete for every fixed $\lambda \ge 2$ [34]. In Section 5 we show that a factor 2 approximation of the path-length of an arbitrary n-vertex graph can be computed in $\mathcal{O}(n^3)$ time. This implies in particular that, for an arbitrary n-vertex graph G, a layout with bandwidth at most $2\,\mathsf{pl}(G)\,\mathsf{bw}(G)$ can be found in $\mathcal{O}(n^3)$ total time.

Additionally, in Section 6 we show that the path-breadth of every permutation graph and every trapezoid graph is 1 and the path-length (and therefore, the path-breadth) of every cocomparability graph and every AT-free graph is at most 2. In Section 7, using some additional structural properties of AT-free graphs, we give a linear time 4-approximation algorithm for the minimum bandwidth problem for AT-free graphs. This result reproduces an approximation result from [31] with a better run-time. Note that the class of AT-free graphs properly contains all permutation graphs, trapezoid graphs and cocomparability graphs; definitions of these graph classes are given in Section 6.

4 Path-length and line-distortion

In this section, we first show that the line-distortion of a graph gives an upper bound on its path-length and then demonstrate that if the path-length of a graph G is bounded by a constant then there is an efficient constant-factor approximation algorithm for the minimum line-distortion problem on G.

4.1 Bound on line-distortion implies bound on path-length

In this subsection we show that the path-length of an arbitrary graph never exceeds its line-distortion. The following inequalities are true.

Proposition 7. *For an arbitrary graph G, $\mathsf{pl}(G) \leq \mathsf{ld}(G)$, $\mathsf{pw}(G) \leq \mathsf{ld}(G)$ and $\mathsf{pb}(G) \leq \lceil \mathsf{ld}(G)/2 \rceil$.*

Proof. It is known (see, e.g., [27]) that every connected graph $G = (V, E)$ has a minimum distortion embedding f into the line ℓ (called a *canonic* embedding) such that $|f(x) - f(y)| = d_G(x, y)$ for every two vertices x, y of G that are placed next to each other in ℓ by f. Assume, in what follows, that f is such a canonic embedding and let $k := \mathsf{ld}(G)$.

Consider the following path-decomposition of G created from f. For each vertex v, form a bag B_v consisting of all vertices of G which are placed by f in the interval $[f(v), f(v) + k]$ of the line ℓ. Order these bags with respect to the left ends of the corresponding intervals. Evidently, for every vertex $v \in V$, $v \in B_v$, i.e., each vertex belongs to a bag. More generally, a vertex u belongs to a bag B_v if and only if $f(v) \leq f(u) \leq f(v) + k$. Since $\mathsf{ld}(G) = k$, for every edge uv of G, $|f(u) - f(v)| \leq k$ holds. Hence, both ends of edge uv belong either to bag B_u (if $f(u) < f(v)$) or to bag B_v (if $f(v) < f(u)$). Now consider three bags B_a, B_b, and B_c with $f(a) < f(b) < f(c)$ and a vertex v of G that belongs to B_a and B_c. We have $f(a) < f(b) < f(c) \leq f(v) \leq f(a) + k < f(b) + k$. Hence, necessarily, v belongs to B_b as well.

It remains to show that each bag B_v, $v \in V$, has in G diameter at most k, radius at most $\lceil k/2 \rceil$ and cardinality at most $k + 1$. Indeed, for any two vertices $x, y \in B_v$, we have $|f(x) - f(y)| \leq k$, i.e., $d_G(x, y) \leq |f(x) - f(y)| \leq k$. Furthermore, any interval $[f(v), f(v) + k]$ (of length k) can have at most $k + 1$ vertices of G as the distance between any two vertices placed by f to this interval is at least 1 ($|f(x) - f(y)| \geq d_G(x, y) \geq 1$). Thus, $|B_v| \leq k + 1$ for every $v \in V$.

Now consider the point $p_v := f(v) + \lfloor k/2 \rfloor$ in the interval $[f(v), f(v) + k]$ of ℓ. Assume, without loss of generality, that p_v is between $f(x)$ and $f(y)$, the images of two vertices x and y of G placed next to each other in ℓ by f. Let $f(x) \leq p_v < f(y)$ (see Fig. 2 for an illustration). Since f is a canonic embedding of G, there must exist a vertex c on a shortest path between x and y such that $d_G(x, c) = p_v - f(x)$ and $d_G(c, y) = f(y) - p_v = d_G(x, y) - d_G(x, c)$. We claim that for every vertex $w \in B_v$, $d_G(c, w) \leq \lceil k/2 \rceil$ holds. Assume $f(w) \geq f(y)$ (the case when $f(w) \leq f(x)$ is similar). Then, we have $d_G(c, w) \leq d_G(c, y) + d_G(y, w) \leq (f(y) - p_v) + (f(w) - f(y)) = f(w) - p_v = f(w) - f(v) - \lfloor k/2 \rfloor \leq k - \lfloor k/2 \rfloor \leq \lceil k/2 \rceil$. □

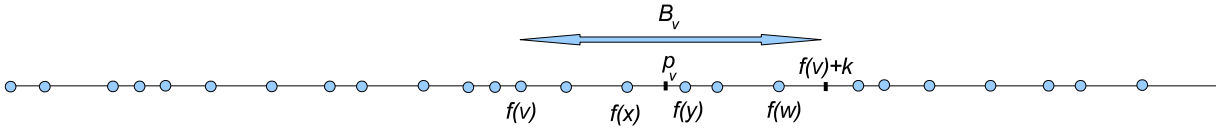

Fig. 2. Illustration to the proof of Proposition 7.

It should be noted that the difference between the path-length and the line-distortion of a graph can be very large. A complete graph K_n on n vertices has path-length 1, whereas the line-distortion of K_n is $n-1$. Note also that the bandwidth and the path-length of a graph do not bound each other. The bandwidth of K_n is $n-1$ while its path-length is 1. On the other hand, the path-length of cycle C_{2n} is n while its bandwidth is 2.

4.2 Line-distortion of graphs with bounded path-length

In this subsection we show that there is an efficient algorithm that for any graph G with $\mathsf{pl}(G) = \lambda$ produces an embedding f of G into the line ℓ with distortion at most $(12\,\lambda + 7)\,\mathsf{ld}(G)$. Again, this statement is true even for all graphs with λ-dominating shortest paths.

We will need the following auxiliary lemma from [4]. We reformulate it slightly. Recall that a subset of vertices of a graph is called *connected* if the subgraph induced by those vertices is connected.

Lemma 2 ([4]). *Any connected subset $S \subseteq V$ of a graph $G = (V, E)$ can be embedded into the line with distortion at most $2|S| - 1$ in time $\mathcal{O}(|V| + |E|)$. In particular, there is a mapping f, computable in $\mathcal{O}(|V| + |E|)$ time, of the vertices from S into points of the line such that $d_G(x, y) \leq |f(x) - f(y)| \leq 2|S| - 1$ for every $x, y \in S$.*

The main result of this subsection is the following.

Proposition 8. *Every graph G with a k-dominating shortest path admits an embedding f of G into the line with distortion at most $(8k + 4)\,\mathsf{ld}(G) + (2k)^2 + 2k + 1$. If a k-dominating shortest path of G is given in advance, then such an embedding f can be found in linear time.*

Proof. Like in the proof of Proposition 5, consider a k-dominating shortest path $P = (x_0, x_1, \ldots, x_i, \ldots, x_j, \ldots, x_q)$ of G and identify by $BFS(P, G)$ the sets X_i, $i \in \{1, \ldots, q\}$. We had $d_G(v, x_i) \leq k$ for every $i \in \{1, \ldots, q\}$ and every $v \in X_i$. It is clear also that each X_i is a connected subset of G. Similar to [4], we define an embedding f of G into the line ℓ by placing all the vertices of X_i before all vertices of X_j, if $i < j$, and by placing vertices within each X_i in accordance with the embedding mentioned in Lemma 2. Also, for each $i \in \{1, \ldots, q-1\}$, leave a space of length $2k + 1$ between the interval of ℓ spanning the vertices of X_i and the interval spanning the vertices of X_{i+1}. See Fig. 3 for an illustration.

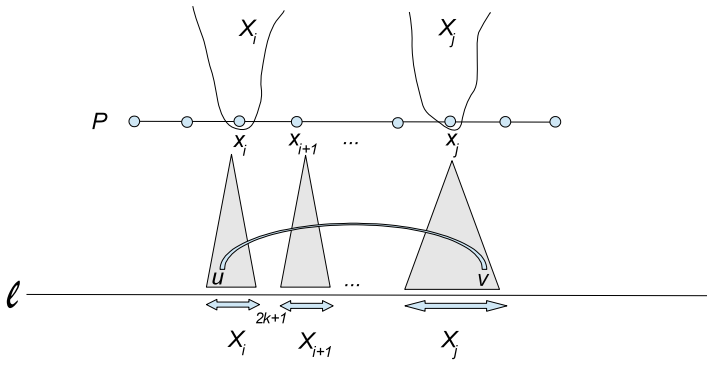

Fig. 3. Illustration to the proof of Proposition 8.

We claim that f is a (non-contractive) embedding with distortion at most $(8k+4)\,\mathsf{ld}(G)+(2k)^2+2k+1$. It is sufficient to show that $d_G(x,y) \leq |f(x)-f(y)|$ for every two vertices of G that are placed next to each other in ℓ by f and that $|f(v)-f(u)| \leq (8k+4)\,\mathsf{ld}(G)+(2k)^2+2k+1$ for every edge uv of G (see, e.g., [4,27]).

From Lemma 2, we know that $d_G(x,y) \leq |f(x)-f(y)| \leq 2|X_h|-1$ for every $x,y \in X_h$ and $h \in \{1,2,\ldots,q\}$. Additionally, for every $x \in X_i$ and $y \in X_{i+1}$ ($i \in \{1,2,\ldots,q-1\}$), we have $d_G(x,y) \leq d_G(x,x_i)+1+d_G(y,x_{i+1}) \leq 2k+1 \leq |f(y)-f(x)|$ (as a space of length $2k+1$ is left between the interval of ℓ spanning the vertices of X_i and the interval spanning the vertices of X_{i+1}). Hence, f is non-contractive.

Consider now an arbitrary edge uv of G and assume $u \in X_i$ and $v \in X_j$ ($i \leq j$). For this edge uv (by Lemma 2) we have $f(v)-f(u) \leq \sum_{h=i}^{j}(2|X_h|-1+2k+1)-2k-1 = 2|\bigcup_{h=i}^{j} X_h|+2k(j-i+1)-2k-1 = 2|\bigcup_{h=i}^{j} X_h|+2k(j-i)-1$. Recall that $d_P(x_i,x_j) = j-i \leq 2k+1$, since P is a shortest path of G and $d_P(x_i,x_j) = d_G(x_i,x_j) \leq d_G(x_i,u)+1+d_G(x_j,v) \leq 2k+1$. Hence, $f(v)-f(u) \leq 2|\bigcup_{h=i}^{j} X_h|+2k(2k+1)-1$.

As in the proof of Proposition 5, $|\bigcup_{h=i}^{j} X_h|-1 \leq (4k+2)\,\mathsf{bw}(G)$. As $\mathsf{bw}(G) \leq \mathsf{ld}(G)$ for every graph G (see, e.g., [27]), we get $f(v)-f(u) \leq 2|\bigcup_{h=i}^{j} X_h|+2k(2k+1)-1 \leq 2(4k+2)\,\mathsf{bw}(G)+2k(2k+1)+1 \leq (8k+4)\,\mathsf{ld}(G)+2k(2k+1)+1$. □

Proposition 8, Corollary 2 and Proposition 2 imply.

Corollary 6. *For every n-vertex m-edge graph G, an embedding into the line with distortion at most $(12\,\mathsf{pl}(G)+7)\,\mathsf{ld}(G)$ can be found in $\mathcal{O}(n^2m)$ time and with distortion at most $(24\,\mathsf{pl}(G)+7)\,\mathsf{ld}(G)$ can be found in $\mathcal{O}(n+m)$ time.*

Proof. See the proof of Corollary 5 and note that, by Proposition 7, $\mathsf{pl}(G) \leq \mathsf{ld}(G)$. Hence, the distortion established in Proposition 8 becomes $\leq (8\,\mathsf{pl}(G)+4)\,\mathsf{ld}(G)+2(2\,\mathsf{pl}(G)+1)\,\mathsf{ld}(G)+1 \leq (12\,\mathsf{pl}(G)+7)\,\mathsf{ld}(G)$, if we use a $\mathsf{pl}(G)$-dominating shortest path, and becomes $\leq (24\,\mathsf{pl}(G)+7)\,\mathsf{ld}(G)$, if we use a $2\,\mathsf{pl}(G)$-dominating shortest path. Algorithm 5 covers both cases. Its runtime and approximation ratio depend on the algorithm to calculate a k-dominating shortest path. □

Algorithm 5: An $\mathcal{O}(\lambda)$-approximation algorithm for computing the minimum line-distortion of a graph G with $\mathsf{pl}(G) \leq \lambda$.

Input: A graph G.
Output: An embedding f of G into the line ℓ.

1 Find a k-dominating shortest path $P = (x_0, x_1, \ldots, x_q)$ using Algorithm 1 or Algorithm 2.
2 Partition V into sets $X_0, X_1, \ldots, X_q$ using a $BFS(P,G)$-tree of G (see the proof of Proposition 8).
3 Create an embedding f of G into the line ℓ by placing all the vertices of X_i before all vertices of X_j, if $i < j$, and by placing vertices within each X_i in accordance with the embedding mentioned in Lemma 2. Also, for each $i \in \{1,\ldots,q-1\}$, leave a space of length $2k+1$ between the interval of ℓ spanning the vertices of X_i and the interval spanning the vertices of X_{i+1}.
4 Output f.

Thus, we have the following interesting conclusion.

Theorem 2. *For every class of graphs with path-length bounded by a constant, there is an efficient constant-factor approximation algorithm for the minimum line-distortion problem.*

Using inequality $\mathsf{pl}(G) \leq \mathsf{ld}(G)$ in Corollary 6 once more, we reproduce a result of [4].

Corollary 7 ([4]). *For every graph G with $\mathsf{ld}(G) = c$, an embedding into the line with distortion at most $\mathcal{O}(c^2)$ can be found in polynomial time.*

It should be noted that, since the difference between the path-length and the line-distortion of a graph can be very large (close to n), the result in Corollary 6 seems to be stronger.

Theorem 1 and Theorem 2 stress the importance of investigating the questions (i) for which particular graph classes there is a constant bound on the path-length and of (ii) how fast can the path-length of an arbitrary graph be computed or sharply estimated. In the next two sections we address some of those questions.

5 Constant-factor approximations of path-length and path-breadth

Let $G = (V, E)$ be an arbitrary graph and let s be an arbitrary vertex of G. A *layering* $\mathcal{L}(s, G)$ of G with respect to a start vertex s is the decomposition of V into layers $L_i = \{u \in V : d_G(s, u) = i\}$, $i = 0, 1, \ldots, q$. For an integer $i \geq 1$ and a vertex $v \in L_i$ denote by $N_G^{\downarrow}(v) = N_G(v) \cap L_{i-1}$ the neighborhood of v in the previous layer L_{i-1}. We can get a path-decomposition of G by adding to each layer L_i ($i > 0$) all vertices from layer L_{i-1} that have a neighbor in L_i, in particular, let $L_i^+ := L_i \cup (\bigcup_{v \in L_i} N_G^{\downarrow}(v))$. Clearly, the sequence $\{L_1^+, \ldots, L_q^+\}$ is a path-decomposition of G and can be constructed in $\mathcal{O}(|E|)$ total time. We call this path-decomposition an *extended layering* of G and denote it by $\mathcal{L}^+(s, G)$.

As shown in the next theorem, this type of path-decomposition has length at most twice as large as the path-length of the graph.

Theorem 3. *For every graph G with $\mathsf{pl}(G) = \lambda$ there is a vertex s such that the length of the extended layering $\mathcal{L}^+(s, G)$ of G is at most $2\,\lambda$. In particular, a factor 2 approximation of the path-length of an arbitrary n-vertex graph can be computed in $\mathcal{O}(n^3)$ total time.*

Proof. Consider a path-decomposition $\mathcal{P}(G) = \{X_1, X_2, \ldots, X_p\}$ of length $\mathsf{pl}(G) = \lambda$ of G. Let s be an arbitrary vertex from X_1. Consider the layering $\mathcal{L}(s, G)$ of G with respect to s where $L_i = \{u \in V : d_G(s, u) = i\}$, $(i = 0, 1, \ldots, q)$. Let x and y be two arbitrary vertices from L_i ($i \in \{1, \ldots, q\}$) and x' and y' be arbitrary vertices from L_{i-1} with $xx', yy' \in E$. We will show that $\max\{d_G(x, y), d_G(x, y'), d_G(x', y)\} \leq 2\,\lambda$. By induction on i, we may assume that $d_G(y', x') \leq 2\,\lambda$ as $x', y' \in L_{i-1}$.

If there is a bag in $\mathcal{P}(G)$ containing both vertices x and y, then $d_G(x, y) \leq \lambda$ and therefore $d_G(x, y') \leq \lambda + 1 \leq 2\,\lambda$, $d_G(y, x') \leq \lambda + 1 \leq 2\,\lambda$. Assume now that all bags containing x are earlier in $\mathcal{P}(G) = \{X_1, X_2, \ldots, X_p\}$ than the bags containing y. Let B be a bag of $\mathcal{P}(G)$ containing both ends of edge xx' (such a bag necessarily exists by properties of path-decompositions). By the position of this bag B in $\mathcal{P}(G)$ and the fact that $s \in X_1$, any shortest path connecting s with y must have a vertex in B. Let w be a vertex of B that is on a shortest path of G connecting vertices s and y and containing edge yy'. Such a shortest path must exist because of the structure of the layering $\mathcal{L}(s, G)$ that starts at s and puts y' and y in consecutive layers. Since $x, x', w \in B$ we have $\max\{d_G(x, w), d_G(x', w)\} \leq \lambda$. If $w = y'$ then we are done; $\max\{d_G(x, y), d_G(x, y'), d_G(x', y)\} \leq \lambda + 1 \leq 2\,\lambda$. So, assume that $w \neq y'$. Since $d_G(x, s) = d_G(s, y) = i$ (by the layering) and $d_G(x, w) \leq \lambda$, we must have $d_G(w, y') + 1 = d_G(w, y) = d_G(s, y) - d_G(s, w) = d_G(s, x) - d_G(s, w) \leq d_G(w, x) \leq \lambda$.

Hence, $d_G(y,x) \leq d_G(y,w) + d_G(w,x) \leq 2\,\lambda$, $d_G(y,x') \leq d_G(y,w) + d_G(w,x') \leq 2\,\lambda$ and $d_G(y',x) \leq d_G(y',w) + d_G(w,x) \leq 2\,\lambda - 1$.

We conclude that the distance between any two vertices of L_i^+ is at most $2\,\lambda$, that is, the length of the tree decomposition $\mathcal{L}^+(s,G)$ of G is at most $2\,\lambda$. □

Algorithm 6 formalizes the method described above.

Algorithm 6: A 2-approximation algorithm for computing the path-length of a graph.

Input: A graph $G = (V,E)$.
Output: A path-decomposition for G.

1 Calculate distances $d_G(u,v)$ for all vertices $u, v \in V$.
2 **foreach** $s \in V$ **do**
3 Calculate a decomposition $\mathcal{L}^+(s,G) = \{L_0^+(s), L_1^+(s), \ldots\}$ with $L_i(s) = \{v \in V : d_G(s,v) = i\}$ and $L_i^+(s) = L_i(s) \cup \{v \in L_{i-1}(s) : N_G(v) \cap L_i(s) \neq \emptyset\}$.
4 Determine the length $l(s)$ of $\mathcal{L}^+(s,G)$
5 Output a decomposition $\mathcal{L}^+(s,G)$ for which $l(s)$ is minimal.

Theorem 4. *For every graph G with $\mathsf{pb}(G) = \rho$ there is a vertex s such that the breadth of the extended layering $\mathcal{L}^+(s,G)$ of G is at most 3ρ. In particular, a factor 3 approximation of the path-breadth of an arbitrary n-vertex graph can be computed in $\mathcal{O}(n^3)$ total time.*

Proof. Since $\mathsf{pl}(G) \leq 2\,\mathsf{pb}(G)$, by Theorem 3, there is a vertex s in G such that the length of extended layering $\mathcal{L}^+(s,G) = \{L_1^+, \ldots, L_q^+\}$ of G is at most 4ρ. Consider a bag L_i^+ of $\mathcal{L}^+(s,G)$ and a family $\mathcal{F} = \{D_G(x,2\rho) : x \in L_i^+\}$ of disks of G. Since $d_G(u,v) \leq 4\rho$ for every pair $u, v \in L_i^+$, the disks of $\mathcal{F}$ pairwise intersect. Hence, by Corollary 4, the disks $\{D_G(x,3\rho) : x \in L_i^+\}$ have a nonempty common intersection. A vertex w from that common intersection has all vertices of L_i^+ within distance at most 3ρ. That is, for each $i \in \{1, \ldots, q\}$ there is a vertex w_i with $L_i^+ \subseteq D_G(w_i, 3\rho)$. □

Combining Theorem 3 and Proposition 6, we obtain the following result.

Theorem 5. *For every n-vertex graph G, a layout f with bandwidth at most $2\,\mathsf{pl}(G)\,\mathsf{bw}(G)$ can be found in $\mathcal{O}(n^3)$ total time.*

6 Bounds on path-length and path-breadth for special graph classes

The class of AT-free graphs contains many intersection families of graphs, among them interval graphs, permutation graphs, trapezoid graphs and cocomparability graphs. These three families of graphs can be defined as follows [5,22]. Consider a line in the plane and n intervals on this line. The intersection graph of such a set of intervals is called an *interval graph*. Consider two parallel lines (upper and lower) in the plane. Assume that each line contains n points, labeled 1 to n. Each two points with the same label define a segment with that label. The intersection graph of such a set of segments between two parallel lines is called a *permutation graph*. Assume now that each of the two parallel lines contains n intervals, labeled 1 to n, and each two intervals with the same label define a trapezoid with that label (a trapezoid can degenerate to a triangle or to a segment). The intersection graph of such a set of trapezoids between two parallel lines is called a *trapezoid graph*. Clearly, every permutation graph is a trapezoid graph, but not vice versa. The class of cocomparability graphs

(which contains all interval graphs and all trapezoid graphs as subclasses) can be defined as the intersection graphs of continuous function diagrams [23], but for this paper it is more convenient to define them via the existence of a special vertex ordering. A graph G is a *cocomparability graph* if it admits a vertex ordering $\sigma = [v_1, v_2, \ldots, v_n]$, called a *cocomparability ordering*, such that for any $i < j < k$, if v_i is adjacent to v_k then v_j must be adjacent to at least one of v_i, v_k. According to [14], such an ordering of a cocomparability graph can be constructed in linear time. Note also that, given a permutation graph G, a *permutation model* (i.e., a set of segments between two parallel lines, defining G) can be found in linear time [14]; a *trapezoid model* for an n-vertex trapezoid graph can be found in $\mathcal{O}(n^2)$ time [35].

In this section we show that the path-breadth of every permutation graph and every trapezoid graph is 1 and the path-length (and therefore, the path-breadth) of every cocomparability graph and every AT-free graph is bounded by 2.

Proposition 9. *If G is a permutation graph, then $\mathsf{pb}(G) = 1$ and, therefore, $\mathsf{pl}(G) \leq 2$. Furthermore, a path-decomposition of G with breadth 1 can be computed in linear time.*

Proof. We assume that a permutation model of G is given in advance (if not, we can compute one for G in linear time [14]). That is, each vertex v of G is associated with a segment $s(v)$ such that $uv \in E$ if and only if segments $s(v)$ and $s(u)$ intersect. In what follows, "u.p." and "l.p." refer to a vertex's point on the upper and lower, respectively, line of the permutation model.

First we compute an (inclusion) maximal independent set M of G in linear time as follows. Put in M (which is initially empty) a vertex x_1 whose u.p. is leftmost. For each $i \geq 2$, select a vertex x_i whose u.p. is leftmost among all vertices whose segments do not intersect $s(x_1), \ldots, s(x_{i-1})$ (in fact, it is enough to check intersection with $s(x_{i-1})$ only). If such a vertex exists, put it in M and continue. If no such vertex exists, $M = \{x_1, \ldots, x_k\}$ has been constructed.

Now, we claim that $\{N_G[x_1], \ldots, N_G[x_k]\}$ is a path-decomposition of G with breadth 1 and, hence, with length at most 2. Clearly, each vertex of G is in some bag since every vertex not in M is adjacent to a vertex in M, by the maximality of M. Consider an arbitrary edge uv of G. Assume that neither u nor v is in M and that the u.p. of u is to the left of the u.p. of v. Necessarily, the l.p. of v is to the left of the l.p. of u, as segments $s(v)$ and $s(u)$ intersect. Assume that the u.p. of u is between the u.p.s of x_i and x_{i+1}. From the construction of M, $s(u)$ and $s(x_i)$ must intersect, i.e., the l.p. of u is to the left of the l.p. of x_i. But then, since the l.p. of v is to the left of the l.p. of x_i, segments $s(v)$ and $s(x_i)$ must intersect, too. Thus, edge uv is in bag $N_G[x_i]$.

To show that all bags containing any particular vertex form a contiguous subsequence of the sequence $N_G[x_1], \ldots, N_G[x_k]$, consider an arbitrary vertex v of G and let $v \in N_G[x_i] \cap N_G[x_j]$ for $i < j$. Consider an arbitrary bag $N_G[x_l]$ with $i < l < j$. We know that vertices $x_i, x_l, x_j \in M$ are pairwise non-adjacent. Furthermore, segment $s(v)$ intersects segments $s(x_i)$ and $s(x_j)$. As segment $s(x_l)$ is between $s(x_i)$ and $s(x_j)$, necessarily, $s(v)$ intersects $s(x_l)$ as well. □

Proposition 10. *If G is an n-vertex trapezoid graph, then $\mathsf{pb}(G) = 1$ and, therefore, $\mathsf{pl}(G) \leq 2$. Furthermore, a path-decomposition of G with breadth 1 can be computed in $\mathcal{O}(n^2)$ time.*

Proof. We will show that every trapezoid graph G is a minor of a permutation graph. Recall that a graph G is called a *minor* of a graph H if G can be formed from H by deleting edges and vertices and by contracting edges.

First, we compute in $\mathcal{O}(n^2)$ time a trapezoid model for G [35]. Then, we replace each trapezoid $\mathcal{T}_i$ in this model with its two diagonals obtaining a permutation model with $2n$ vertices. Let H be the permutation graph of this permutation model. It is easy to see that two trapezoids $\mathcal{T}_1$ and $\mathcal{T}_2$ intersect if and only if a diagonal of $\mathcal{T}_1$ and a diagonal of $\mathcal{T}_2$ intersect.

Now, G can be obtained back from H by a series of n edge contractions; for each trapezoid $\mathcal{T}_i$, contract the edge of H that corresponds to two diagonals of $\mathcal{T}_i$.

Since contracting edges does not increase the path-breadth (see [11]), we get $\mathsf{pb}(G) = \mathsf{pb}(H) = 1$ by Proposition 9. Any path-decomposition of H with breadth 1 is a path-decomposition of G with breadth 1. □

For the proof of the next result we will need a special vertex ordering $\sigma : V \to \{1, \ldots, n\}$ produced by a so-called *Lexicographic-Breadth-First-Search* (LBFS for short) which is a refinement of a standard *Breadth-First-Search* (BFS). LBFS(s, G) starts at some start vertex s, orders the vertices of a graph G by assigning numbers from n to 1 to the vertices in the order as they are discovered by the following search process. Each vertex v has a *label* consisting of a (revers) ordered list of the *numbers* of those neighbors of v that were already visited by the LBFS; initially this label is empty. LBFS starts with some vertex s, assigns number n to s, and adds this number to the end of the label of all un-numbered neighbors of s. Then, in each step, LBFS selects the un-numbered vertex v with the lexicographically largest label, assigns the next available number k to v, and adds this number to the end of the labels of all un-numbered neighbors of v. An ordering σ of the vertex set of a graph generated by LBFS(s, G) is called a LBFS(s, G)-ordering. Note that the closer a vertex is to s in G the larger its number is in σ. It is known that a LBFS-ordering of an arbitrary graph can be generated in linear time [22].

Proposition 11. *If G is an n-vertex AT-free graph, then $\mathsf{pb}(G) \leq \mathsf{pl}(G) \leq 2$. Furthermore, a path-decomposition of G with length at most 2 can be computed in $\mathcal{O}(n^2)$ time.*

Proof. Let s be an arbitrary vertex of G and x be a vertex last visited (numbered 1) by an LBFS(s, G). Let σ be an LBFS(x, G)-ordering of vertices of G. Clearly, σ can be generated in linear time as one needs only 2 scans of LBFS to do that. The following useful result was proven in [8].

Claim 1: [8] For every vertex y of an AT-free graph G, the pair x, y is a 1-dominating pair of the subgraph $G_{\geq\sigma(y)}$ of G induced by vertices $\{z \in V : \sigma(y) \leq \sigma(z) \leq \sigma(x) = n\}$. In particular, x and the vertex last visited by LBFS(x, G) constitute a 1-dominating pair of G.

Let now $\mathcal{L}(u, G) = \{L_0, \ldots, L_k\}$ with $L_i = \{u \in V : d_G(u, x) = i\}$ be a layering of G produced by LBFS(x, G).

Claim 2: For every integer $i \geq 1$ and every two non-adjacent vertices $u, v \in L_i$ of an AT-free graph G, $\sigma(v) < \sigma(u)$ implies $N^{\downarrow}_G(v) \subseteq N^{\downarrow}_G(u)$. In particular, $d_G(v, u) \leq 2$ holds for every $u, v \in L_i$ and every i.

Proof (of Claim 2). Consider an arbitrary neighbor $w \in L_{i-1}$ of v and a shortest path P from v to x in G containing w. Since $\sigma(v) < \sigma(u)$, by Claim 1, path P must dominate vertex u. Since u and v are not adjacent, u is in L_i and all vertices of $P \setminus \{v, w\}$ belong to layers L_j with $j < i - 1$, vertex u must be adjacent to w. □(Claim)

We can transform an AT-free graph $G = (V, E)$ into an interval graph $G^+ = (V, E^+)$ by applying the following two operations:

(1) (*make layers complete graphs*) In each layer L_i, make every two vertices $u, v \in L_i$ adjacent to each other in G^+;
(2) (*make down-neighborhoods of adjacent vertices of a layer comparable, too*) For each i and every edge uv of G with $u, v \in L_i$ and $\sigma(v) < \sigma(u)$, make every $w \in N^{\downarrow}_G(v)$ adjacent to u in G^+.

Clearly, for every edge uw of G^+ added by operation (2), $d_G(u,w) \leq 2$ holds. Also, for every edge uv of G^+ added by operation (1), $d_G(u,v) \leq 2$ holds by Claim 2. Thus, we have.

Claim 3: G^+ is a subgraph of G^2.

Next we show that G^+ is an interval graph.

Claim 4: G^+ is an interval graph.

Proof (of Claim 4). It is known [37] that a graph is an interval graph if and only if its vertices admit an *interval ordering*, i.e., an ordering $\tau : V \to \{1, \ldots, n\}$ such that for every choice of vertices a, b, c with $\tau(a) < \tau(b) < \tau(c)$, $ac \in E$ implies $bc \in E$. We show here that the LBFS(x, G)-ordering σ of G is an interval ordering of G^+. Recall that, for every $v \in L_i$ and $u \in L_j$ with $i > j$, it holds that $\sigma(v) < \sigma(u)$ (as σ is an LBFS-ordering). Consider three arbitrary vertices a, b, c of G and assume that $\sigma(a) < \sigma(b) < \sigma(c)$ and $ac \in E^+$. Assume also that $a \in L_i$ for some i. If c belongs to L_i, then b must be in L_i as well and, hence, $bc \in E^+$ due to operation (1). If both b and c are in L_{i-1}, then again $bc \in E^+$ due to operation (1). Consider now the remaining case, i.e., $a, b \in L_i$ and $c \in L_{i-1}$. If $ac \in E$ then $bc \in E^+$ because either $ab \in E$ and thus operation (2) applies, or $ab \notin E$ and thus Claim 2 implies $bc \in E^+$. If $ac \in E^+ \setminus E$ then, according to operation (2), edge ac was created in G^+ because some vertex $a' \in L_i$ existed such that $\sigma(a') < \sigma(a)$ and $a'a, a'c \in E$. Since $a'c \in E$ and $\sigma(a') < \sigma(b) < \sigma(c)$, as before, $bc \in E^+$ must hold. □(Claim)

To complete the proof of Proposition 11, we recall that a graph is an interval graph if and only if it has a path-decomposition with each bag being a maximal clique (see, e.g., [9,19,20,22]). Furthermore, such a path-decomposition of an interval graph can easily be computed in linear time. Let $\mathcal{P}(G^+) = \{X_1, X_2, \ldots, X_q\}$ be a path-decomposition of our interval graph G^+. Then, $\mathcal{P}(G) := \mathcal{P}(G^+) = \{X_1, X_2, \ldots, X_q\}$ is a path decomposition of G with length at most 2 since, for every edge uv of G^+, the distance in G between u and v is at most 2, as shown in Claim 3. □

Algorithm 7 formalizes the steps described in the previous proof.

Algorithm 7: Computing a path-decomposition of length at most 2 for a given AT-free graph.

Input: An AT-free graph $G = (V, E)$.
Output: A path-decomposition of G.

1 Calculate an LBFS(s, G) ordering σ with an arbitrary start vertex $s \in V$. Let x be the last visited vertex, i.e., $\sigma(x) = 1$.
2 Calculate an LBFS(x, G) ordering σ'.
3 Set $E^+ := E$.
4 **foreach** *vertex pair u, v with $d_G(x,u) = d_G(x,v)$ and $\sigma'(u) < \sigma'(v)$* **do**
5 Add uv to E^+.
6 For each $w \in N_G(u)$ with $\sigma'(v) < \sigma'(w)$, add vw to E^+.
7 Calculate a path-decomposition $\mathcal{P}(G^+)$ of the interval graph $G^+ = (V, E^+)$ by determining the maximal cliques of G^+.
8 Output $\mathcal{P}(G^+)$.

As the class of cocomparability graphs is a proper subclass of AT-free graphs, we obtain the following corollary.

Corollary 8. *If G is an n-vertex cocomparability graph, then $\mathsf{pb}(G) \leq \mathsf{pl}(G) \leq 2$. Furthermore, a path-decomposition of G with length at most 2 can be computed in $\mathcal{O}(n^2)$ time.*

The complement of an induced cycle on six vertices shows that the bound 2 on the path-breadth of cocomparability graphs (and therefore, of AT-free graphs) is sharp. Indeed, the edge set of $\overline{C}_6$ forms a bramble but no vertex 1-dominates all edges, implying, by Proposition 4, that $\mathsf{pb}(\overline{C}_6) = 2$.

Since the minimum line-distortion problem is NP-hard on cocomparability graphs [26], it is NP-hard also on bounded path-length graphs.

Corollary 9. *The minimum line-distortion problem is NP-hard on bounded path-length graphs.*

We know that the minimum bandwidth problem is NP-hard even on bounded path-width graphs (e.g., even on caterpillars of hair-length at most 3 [36,13]). Recently, in [41], it was shown that the minimum bandwidth problem is NP-hard also on so-called *convex bipartite graphs*. A bipartite graph $G = (U, V; E)$ is said to be *convex* if for one of its parts, say U, there is an ordering $(u_1, u_2, \ldots, u_q)$ such that for all $v \in V$ the vertices adjacent to v are consecutive. It is easy to see that, in this case, $\{N_G[u_1], N_G[u_2], \ldots, N_G[u_q]\}$ is a path-decomposition of G of breadth 1. Note that, given a convex bipartite graph $G = (U, V; E)$, a proper ordering $(u_1, u_2, \ldots, u_q)$ of U can be found in linear time [6]. Thus, the following two results are true.

Proposition 12. *If G is a convex bipartite graph, then $\mathsf{pb}(G) = 1$ and, therefore, $\mathsf{pl}(G) \leq 2$. Furthermore, a path-decomposition of G with breadth 1 can be computed in linear time.*

Corollary 10. *The minimum bandwidth problem is NP-hard on bounded path-length graphs.*

7 Constant-factor approximation of line-distortions of AT-free graphs

From Theorem 2 and results of the previous section, it follows already that there is an efficient constant-factor approximation algorithm for the minimum line-distortion problem on such particular graph classes as permutation graphs, trapezoid graphs, cocomparability graphs as well as AT-free graphs. Recall that for arbitrary (unweighted) graphs the minimum line-distortion problem is hard to approximate within a constant factor [4]. Furthermore, the problem remains NP-hard even when the input graph is restricted to a chordal, cocomparability, or AT-free graph [26]. Polynomial-time constant-factor approximation algorithms were known only for split and cocomparability graphs; [26] gave efficient 6-approximation algorithms for both graph classes. As far as we know, for AT-free graphs (the class which contains all cocomparability graphs), no prior efficient approximation algorithm was known.

In this section, using additional structural properties of AT-free graphs and ideas from Section 4.2, we give a better approximation algorithm for all AT-free graphs; more precisely, we give an 8-approximation algorithm that runs in linear time.

The following nice structural result from [31] will be very useful.

Lemma 3 ([31]). *Let $G = (V, E)$ be an AT-free graph. Then, there is a dominating path $\pi = (v_0, \ldots, v_k)$ and a layering $\mathcal{L} = \{L_0, \ldots, L_k\}$ with $L_i = \{u \in V : d_G(u, v_0) = i\}$ such that for all $u \in L_i$ $(i \geq 1)$, $uv_i \in E$ or $uv_{i-1} \in E$. Computing π and $\mathcal{L}$ can be done in linear time.*

Theorem 6. *There is a linear time algorithm to compute an 8-approximation of the line-distortion of an AT-free graph.*

Proof. Let G be an AT-free graph. We first compute a path $\pi = (v_0, \ldots, v_k)$ and a layering $\mathcal{L} = \{L_0, \ldots, L_k\}$ as defined in Lemma 3. To define an embedding f of G into the line, we partition every layer L_i in three sets: $\{v_i\}$, $X_i = \{x : x \in L_i, v_i x \in E\}$, and $\overline{X}_i = L_i \setminus (\{v_i\} \cup X_i)$ (see Fig. 4). Note that if $x \in \overline{X}_i$, then $v_{i-1}x \in E$. Since each vertex in X_i is adjacent to v_i and each vertex in $\overline{X}_i$ is adjacent to v_{i-1}, for all $x, y \in X_i$, $d_G(x, y) \leq 2$, and for all $x, y \in \overline{X}_i$, $d_G(x, y) \leq 2$. Also, for all $x \in X_i$ and $y \in \overline{X}_i$, $d_G(x, y) \leq 3$. The embedding f places vertices of G into the line in the following order: $(v_0, \ldots, v_{i-1}, \overline{X}_i, X_i, v_i, \overline{X}_{i+1}, X_{i+1}, v_{i+1}, \ldots, v_k)$. Between every two vertices x, y placed next to each other on the line, to guarantee non-contractiveness, f leaves a space of length $d_G(x, y)$ (which is either 1 or 2 or 3, where 3 occurs only when $x \in \overline{X}_i$ and $y \in X_i$ for some i).

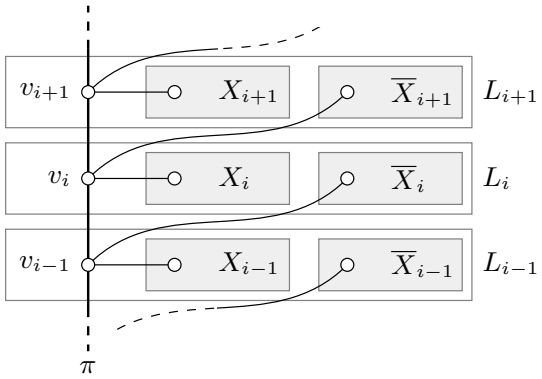

Fig. 4. Layering of an AT-free graph; illustration to the proof of Theorem 6.

Now we will show that f approximates the minimum line-distortion of G. Since $\mathcal{L}$ is a BFS layering started from v_0, i.e., it represents the distances of vertices from v_0, there is no edge xy with $x \in L_{i-1}$ and $y \in L_{i+1}$. Also note that $D_G(v_i, 2) \supseteq L_i \cup L_{i+1} \cup \{v_{i-1}\}$. By the definition of f, for all $xy \in E$ with $x, y \in L_i \cup L_{i+1}$, $|f(x) - f(y)| < |f(v_{i-1}) - f(v_{i+1})|$. Therefore, counting how many vertices are placed by f between $f(v_{i-1})$ and $f(v_{i+1})$ and the distance in G between vertices placed next to each other, we get $|f(x) - f(y)| \leq 2(|D_G(v_i, 2)| - 2) + 2 = 2(|D_G(v_i, 2)| - 1)$. Using Lemma 1 and the fact that $\mathsf{bw}(G) \leq \mathsf{ld}(G)$, we get $|f(x) - f(y)| \leq 8\,\mathsf{ld}(G)$ for all $xy \in E$. □

Algorithm 8 formalizes the method described above.

Algorithm 8: An 8-approximation algorithm for the minimum line-distortion of an AT-free graph.

Input: An AT-free graph $G = (V, E)$.
Output: An embedding f of G into the line.

1 Compute a path $\pi = (v_0, \ldots, v_k)$ and a layering $\mathcal{L} = \{L_0, \ldots, L_k\}$ as defined in Lemma 3.
2 Partition each layer L_i into three sets: $\{v_i\}$, $X_i = \{x : x \in L_i, v_i x \in E\}$, and $\overline{X}_i = L_i \setminus (\{v_i\} \cup X_i)$.
3 Create an embedding f by placing the vertices of G into the line in the order $(v_0, \ldots, v_{i-1}, \overline{X}_i, X_i, v_i, \overline{X}_{i+1}, X_{i+1}, v_{i+1}, \ldots, v_k)$.
4 Between every two consecutive vertices x, y on the line, leave a space of length $d_G(x, y)$.
5 Output f.

It is easy to see that the order in which the vertices of G are placed by f into the line gives also a layout of G with bandwidth at most $4\,\mathsf{bw}(G)$. This reproduces an approximation result from [31] (in fact, their algorithm has complexity $\mathcal{O}(m + n\log n)$ for an n-vertex m-edge graph, since it involves an $\mathcal{O}(n\log n)$ time algorithm from [1] to find an optimal layout for a caterpillar with hair-length at most 1).

Corollary 11 ([31]). *There is a linear time algorithm to compute a 4-approximation of the minimum bandwidth of an AT-free graph.*

Combining Proposition 11 and Proposition 6, we obtain also the following result from [31] as a corollary.

Corollary 12 ([31]). *There is an $\mathcal{O}(m + n\log^2 n)$ time algorithm to compute a 2-approximation of the minimum bandwidth of an AT-free graph.*

8 Concluding remarks

In this paper we have shown that if a graph G has a k-dominating shortest path, where k is a constant, or the path-length of G is bounded by a constant, then both the minimum line-distortion problem and the minimum bandwidth problem on G can be efficiently approximated within constant factors. As AT-free graphs, cocomparability graphs, permutation graphs, trapezoid graphs, convex bipartite graphs, caterpillars with hairs of bounded length, all have bounded path-length or have k-dominating shortest paths with constant k, they admit constant factor approximations of the minimum bandwidth and the minimum line-distortion. Thus, the constant factor approximation results of [26,31,41] become special cases of our results.

We conclude this paper with a few open questions. We have presented a 2-approximation algorithm for computing the path-length of a general graph but we do not know the complexity status of this problem. So, our first open question is the following.

1) Is it NP-complete to decide whether a graph has path-length at most k ($k > 1$)?

We gave a first constant-factor approximation (8-approximation) algorithm for the minimum line-distortion problem on AT-free graphs and reproduced a 4-approximation and a 2-approximation for the minimum bandwidth on such graphs.

2) Does there exist a better approximation algorithm for the minimum line-distortion problem on AT-free graphs?

We have mentioned that the minimum bandwidth problem is notoriously hard even on bounded path-width graphs (e.g., even on caterpillars of hair-length at most 3 [36,13]) and even on bounded path-length graphs (e.g., even on convex bipartite graphs [41]). Since the minimum line-distortion problem is NP-hard on cocomparability graphs [26], it is NP-hard also on bounded path-length graphs. However, the status of the minimum line-distortion problem on bounded path-width graphs is unknown. See Table 5 for a summary.

3) Is the minimum line-distortion problem on bounded path-width graphs NP-hard?

We are also interested in a more general question.

4) Is there a better hardness result (better than a constant [4]) for the line-distortion problem in general graphs?

Table 5. Hardness results for the minimum line-distortion problem and the minimum bandwidth problem on graphs with bounded path-with or bounded path-length.

	$\mathsf{pw}(G) \leq c$	$\mathsf{pl}(G) \leq c$
bandwidth	NP-hard (caterpillars with hair-length at most 3 [36])	NP-hard (convex bipartite graphs [41])
line-distortion	?	NP-hard (cocomparability graphs [26])

References

1. S.F. Assman, G.W. Peck, M.M. Syslo, J. Zak, The bandwidth of caterpillars with hairs of length 1 and 2, *SIAM J. Alg. Disc. Meth.* 2 (1981), 387–392.
2. G. Blache, M. Karpinski, J.Wirtgen, On approximation intractability of the bandwidth problem, *Technical report TR98-014,* University of Bonn, 1997.
3. M. Bădoiu, J. Chuzhoy, P. Indyk, A. Sidiropoulos, Low-distortion embeddings of general metrics into the line, In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC 2005)*, Baltimore, MD, USA, May 22-24, 2005, ACM, pp. 225–233.
4. M. Bădoiu, K. Dhamdhere, A. Gupta, Y. Rabinovich, H. Raecke, R. Ravi, and A. Sidiropoulos, Approximation algorithms for low-distortion embeddings into low-dimensional spaces, Proceedings of the ACM/SIAM Symposium on Discrete Algorithms, 2005.
5. A. Brandstädt, V.B. Le and J. Spinrad, Graph Classes: A Survey, *SIAM,* Philadelphia, 1999.
6. K.S. Booth, G.S. Lueker, Testing for the consecutive ones property, interval graphs, and graph planarity using PQ-tree algorithms, *J. Comput. Syst. Sci.* 13(3) (1976), 335–379.
7. D. G. Corneil, S. Olariu, L. Stewart, Asteroidal Triple-Free Graphs, *SIAM Journal on Discrete Mathematics* 10 (1997), 399–430.
8. D. G. Corneil, S. Olariu, L. Stewart, Linear Time Algorithms for Dominating Pairs in Asteroidal Triple-free Graphs, *SIAM J. Computing* 28 (1997), 292–302.
9. R. Diestel, *Graph Theory,* second edition, Graduate Texts in Mathematics, vol. 173, Springer, 2000.
10. Y. Dourisboure and C. Gavoille, Tree-decompositions with bags of small diameter, *Discr. Math.* 307 (2007) 208–229.
11. F.F. Dragan, E. Köhler, An Approximation Algorithm for the Tree t-Spanner Problem on Unweighted Graphs via Generalized Chordal Graphs, Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques - Proceedings of the 14th International Workshop, APPROX 2011, and 15th International Workshop, RANDOM 2011, Princeton, NJ, USA, August 17-19, 2011, *Lecture Notes in Computer Science* 6845, Springer, pp. 171-183; *Algorithmica* 69 (2014), 884–905.
12. F.F. Dragan, E. Köhler and A. Leitert, Line-distortion, Bandwidth and Path-length of a graph, In Proceedings of *14th Scandinavian Symposium and Workshops on Algorithm Theory* (SWAT 2014), July 2-4 2014. Copenhagen, Denmark, *Lecture Notes in Computer Science* 8503, 2014, pp. 146-257.
13. Ch. Dubey, U. Feige, W. Unger, Hardness results for approximating the bandwidth, *Journal of Computer and System Sciences* 77 (2011), 62–90.
14. R.M. McConnell and J.P. Spinrad, Linear-time transitive orientation, In *Proceedings of the Eighth Annual ACM-SIAM Symposium on Discrete Algorithms,* New Orleans, Louisiana, 5-7 January 1997, pp. 19-25.
15. U. Feige, Approximating the bandwidth via volume respecting embedding, *J. of Computer and System Science,* 60 (2000), 510–539.
16. U. Feige, K. Talwar, Approximating the Bandwidth of Caterpillars, In Proceedings of *8th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX 2005) and 9th International Workshop on Randomization and Computation (RANDOM 2005)*, Berkeley, CA, USA, August 22-24, 2005, *Lecture Notes in Computer Science* 3624, 2005, pp 62–73.
17. M.R. Fellows, F.V. Fomin, D. Lokshtanov, E. Losievskaja, F.A. Rosamond, S. Saurabh, Distortion Is Fixed Parameter Tractable, In *ICALP* 2009, pp. 463–474.
18. F.V. Fomin, D. Lokshtanov, S. Saurabh, An exact algorithm for minimum distortion embedding, *Theor. Comput. Sci.* 412 (2011), 3530–3536.
19. D.R. Fulkerson, O.A. Gross, Incidence matrices and interval graphs, *Pacific J. Math.* 15 (1965) 835–855.
20. P. C. Gilmore and A. J. Hoffman, A characterization of comparability graphs and interval graphs, *Canad. J. Math.* 16 (1964), pp. 539–548.
21. P.A. Golovach, P. Heggernes, D. Kratsch, D. Lokshtanov, D. Meister, S. Saurabh, Bandwidth on AT-free graphs, *Theor. Comput. Sci.* 412 (2011), 7001–7008.
22. M.C. Golumbic, *Algorithmic Graph Theory and Perfect Graphs*, Academic Press, New York, 1980.

23. M.C. Golumbic, D. Rotem, J. Urrutia, Comparability graphs and intersection graphs, *Discrete Mathematics* 43(1983), 37–46.
24. A. Gupta, Improved Bandwidth Approximation for Trees and Chordal Graphs, *J. Algorithms* 40(2001), 24–36.
25. P. Heggernes, D. Kratsch, D. Meister, Bandwidth of bipartite permutation graphs in polynomial time, *Journal of Discrete Algorithms* 7 (2009), 533–544.
26. P. Heggernes and D. Meister, Hardness and approximation of minimum distortion embeddings, *Information Processing Letters* 110 (2010), 312–316.
27. P. Heggernes, D. Meister, and A. Proskurowski, Computing minimum distortion embeddings into a path of bipartite permutation graphs and threshold graphs, *Theoretical Computer Science* 412 (2011), 1275–1297.
28. P. Indyk, Algorithmic applications of low-distortion geometric embeddings, Proceedings of FOCS 2001, pp. 10–35, IEEE, 2005.
29. P. Indyk and J. Matousek, Low-distortion embeddings of finite metric spaces, Handbook of Discrete and Computational Geometry, second ed., pp. 177–196, CRC press, 2004.
30. D. J. Kleitman, R. V. Vohra, Computing the bandwidth of interval graphs, *SIAM J. Disc. Math.* 3 (1990), 373–375.
31. T. Kloks, D. Kratsch, H. Müller, Approximating the Bandwidth for Asteroidal Triple-Free Graphs, *J. Algorithms* 32 (1999), 41–57.
32. D. Kratsch, J. Spinrad, Between $\mathcal{O}(nm)$ and $\mathcal{O}(n^\alpha)$, In *Proceedings of the fourteenth annual ACM-SIAM symposium on Discrete algorithms (SODA 2003)*, ACM, pp. 709–716.
33. D. Kratsch, L. Stewart, Approximating Bandwidth by Mixing Layouts of Interval Graphs, *SIAM J. Discrete Math.* 15(2002), 435–449.
34. D. Lokshtanov. On the complexity of computing treelength. *Discrete Applied Mathematics*, 158(7):820–827, 2010.
35. T.H. Ma and J.P. Spinrad, On the two-chain subgraph cover and related problems, *J. of Algorithms*, 17 (1994), 251–268.
36. B. Monien, The Bandwidth-Minimization Problem for Caterpillars with Hair Length 3 is NP-Complete, *SIAM J. Alg. Disc. Meth.* 7 (1986), 505–512.
37. S. Olariu, An optimal greedy heuristic to color interval graphs, *Inform. Process. Lett.* 37 (1991), 65–80.
38. A. Proskurowski, J.A. Telle, Classes of graphs with restricted interval models, *Discrete Mathematics and Theoretical Computer Science* 3 (1999), 167-176.
39. H. Räcke, Lecture notes at http://ttic.uchicago.edu/∼ harry/teaching/pdf/lecture15.pdf
40. N. Robertson, P. Seymour, Graph minors. I. Excluding a forest, *Journal of Combinatorial Theory, Series B* 35 (1983), 39–61.
41. A.M.S. Shrestha, S. Tayu, S. Ueno, Bandwidth of convex bipartite graphs and related graphs, *Information Processing Letters* 112 (2012), 411–417.
42. A. P. Sprague, An $\mathcal{O}(n \log n)$ algorithm for bandwidth of interval graphs, *SIAM J. Disc. Math.* 7 (1994), 213–220.
43. J. B. Tenenbaum, V. de Silva, and J. C. Langford, A global geometric framework for nonlinear dimensionality reduction, *Science* 290 (2000), 2319–2323.